

The Constrained Dantzig Selector with Enhanced Consistency *

Yinfei Kong

*Department of Preventive Medicine
University of Southern California
Los Angeles, CA 90089, USA*

YINFEIKO@USC.EDU

Zemin Zheng

*Department of Statistics and Finance
University of Science and Technology of China
Hefei, Anhui 230026, China*

ZHENGZM@USTC.EDU.CN

Jinchi Lv

*Data Sciences and Operations Department
Marshall School of Business
University of Southern California
Los Angeles, CA 90089, USA*

JINCHILV@MARSHALL.USC.EDU

Editor: Sara van de Geer

Abstract

The Dantzig selector has received popularity for many applications such as compressed sensing and sparse modeling, thanks to its computational efficiency as a linear programming problem and its nice sampling properties. Existing results show that it can recover sparse signals mimicking the accuracy of the ideal procedure, up to a logarithmic factor of the dimensionality. Such a factor has been shown to hold for many regularization methods. An important question is whether this factor can be reduced to a logarithmic factor of the sample size in ultra-high dimensions under mild regularity conditions. To provide an affirmative answer, in this paper we suggest the constrained Dantzig selector, which has more flexible constraints and parameter space. We prove that the suggested method can achieve convergence rates within a logarithmic factor of the sample size of the oracle rates and improved sparsity, under a fairly weak assumption on the signal strength. Such improvement is significant in ultra-high dimensions. This method can be implemented efficiently through sequential linear programming. Numerical studies confirm that the sample size needed for a certain level of accuracy in these problems can be much reduced.

Keywords: Sparse Modeling, Compressed Sensing, Ultra-high Dimensionality, Dantzig Selector, Regularization Methods, Finite Sample

*. This work was supported by the NSF CAREER Award DMS-0955316, a grant from the Simons Foundation, and USC Diploma in Innovation Grant. We sincerely thank the Action Editor and two referees for their valuable comments that have helped improve the paper significantly.

1. Introduction

Due to the advances of technologies, big data problems appear increasingly common in the domains of molecular biology, machine learning, and economics. It is appealing to design procedures that can provide a recovery of important signals, among a pool of potentially huge number of signals, to a desired level of accuracy. As a powerful tool for producing interpretable models, sparse modeling via regularization has gained popularity for analyzing large-scale data sets. A common feature of the theories for many regularization methods is that the rates of convergence under the prediction or estimation loss usually involve the logarithmic factor of the dimensionality $\log p$; see, for example, Bickel et al. (2009), van de Geer et al. (2011), and Fan and Lv (2013), among many others. From the asymptotic point of view, such a factor may be negligible or insignificant when p is not too large. It can, however, become no longer negligible or even significant in finite samples with large dimensionality, particularly when a relatively small sample size is considered or preferred. In such cases, the recovered signals and sparse models may tend to be noisy. Another consequence of this effect is that many noise variables tend to appear together with recovered important ones (Candès et al., 2008). An important and interesting question is whether such a factor can be reduced, say, to a logarithmic factor of the sample size $\log n$, from ultra-high dimensions under mild regularity conditions.

In high-dimensional variable selection, there is a fast growing literature for different kinds of regularization methods with well established rates of convergence under the estimation and prediction losses. For instance, Candès and Tao (2007) proposed the L_1 -regularization approach of the Dantzig selector by relaxing the normal equation to allow for correlation between the residual vector and all the variables, which can recover sparse signals as accurately as the ideal procedure, knowing the true underlying sparse model, up to a logarithmic factor $\log p$. Later, Bickel et al. (2009) showed that the popularly used Lasso estimator (Tibshirani, 1996) exhibits similar behavior as the Dantzig selector with a $\log p$ factor in the oracle inequalities for the prediction risk and bounds on the estimation losses in general nonparametric regression models. Furthermore, it has been proved in Raskutti et al. (2011) that the minimax rates of convergence for estimating the true coefficient vector in the high-dimensional linear regression model involve a logarithmic factor $\log p$ in the L_2 -loss and prediction loss, under some regularity conditions on the design matrix. Other work for desired properties such as the oracle property on the L_1 and more general regularization includes Fan and Li (2001), Fan and Peng (2004), Zou and Hastie (2005), Zou (2006), van de Geer (2008), Lv and Fan (2009), Antoniadis et al. (2010), Städler et al. (2010), Fan and Lv (2011), Negahban et al. (2012), Chen et al. (2013), and Fan and Tang (2013), among many others.

A typical way for reducing the logarithmic factor $\log p$ to $\log n$ is through model selection consistency, where the estimator selects exactly the support of the true coefficient vector, that is, the set of variables with nonzero coefficients. We refer the reader to Zhao and Yu (2006), Wainwright (2009), Zhang (2010), and Zhang (2011) for analysis of model selection consistency of regularization methods. Since the true parameters are assumed to be sparse in the high-dimensional setting, consistent variable selection can greatly lessen the analytical complexity from large dimensionality p to around oracle model size. Once model selection consistency is established for the estimator with significant probability, an analysis

constrained on that event will give a factor $\log n$ instead of $\log p$ in the rates of convergence under various estimation and prediction losses. However, to obtain model selection consistency it usually requires a uniform signal strength condition that the minimum magnitude of nonzero coefficients is at least of order $\{s(\log p)/n\}^{1/2}$ (Fan and Lv, 2013 and Zheng et al., 2014), where s stands for the size of the low-dimensional parameter vector. We doubt the necessity of the model selection consistency and the uniform signal strength of order $\{s(\log p)/n\}^{1/2}$ for achieving the logarithmic factor $\log n$ in ultra-high dimensionality.

In this paper, we suggest the constrained Dantzig selector to study the rates of convergence with weaker signal strength assumption. The constrained Dantzig selector replaces the constant Dantzig constraint on correlations between the variables and the residual vector by a more flexible one, and considers a constrained parameter space distinguishing between zero parameters and significantly nonzero parameters. The main contributions of this paper are threefold. First, the convergence rates for the constrained Dantzig selector are shown to be within a factor of $\log n$ of the oracle rates instead of $\log p$, a significant improvement in the case of ultra-high dimensionality and relatively small sample size. It is appealing that such an improvement is made with a fairly weak assumption on the signal strength without requiring model selection consistency. To the best of our knowledge, this assumption seems to be the weakest one in the literature of similar results; see, for example, Bickel et al. (2009) and Zheng et al. (2014). Two parallel theorems, under the uniform uncertainty principle condition and the restricted eigenvalue assumption, are established on the properties of the constrained Dantzig selector for compressed sensing and sparse modeling, respectively. Second, compared to the Dantzig selector, theoretical results of this paper show that the number of falsely discovered signs of our new selector, with an explicit inverse relationship to the signal strength, is controlled as a possibly asymptotically vanishing fraction of the true model size. Third, an active-set based algorithm is introduced to implement the constrained Dantzig selector efficiently. An appealing feature of this algorithm is that its convergence can be checked easily.

The rest of the paper is organized as follows. In Section 2, we introduce the constrained Dantzig selector. We present its compressed sensing and sampling properties in Section 3. In Section 4, we discuss the implementation of the method and present several simulation and real data examples. We provide some discussions of our results and some possible extensions of our method in Section 5. All technical details are relegated to the Appendix.

2. The Constrained Dantzig selector

To simplify the technical presentation, we adopt the model setting in Candès and Tao (2007) and present the main ideas focusing on the linear regression model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_n)^T$ is an n -dimensional response vector, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_p)$ is an $n \times p$ design matrix consisting of p covariate vectors \mathbf{x}_j 's, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a p -dimensional regression coefficient vector, and $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)^T \sim N(\mathbf{0}, \sigma^2 I_n)$ for some positive constant σ is an n -dimensional error vector independent of \mathbf{X} . The normality assumption is considered for simplicity, and all the results in the paper can be extended to the cases of bounded errors

or light-tailed error distributions without much difficulty. See, for example, the technical analysis in Fan and Lv (2011) and Fan and Lv (2013).

In both problems of compressed sensing and sparse modeling, we are interested in recovering the support and nonzero components of the true regression coefficient vector $\beta_0 = (\beta_{0,1}, \dots, \beta_{0,p})^T$, which we assume to be sparse with s nonzero components, for the case when the dimensionality p may greatly exceed the sample size n . Throughout this paper, p is implicitly understood as $\max\{n, p\}$ and $s \leq \min\{n, p\}$ to ensure model identifiability. To align the scale of all covariates, we assume that each column of X , that is, each covariate vector \mathbf{x}_j , is rescaled to have L_2 -norm $n^{1/2}$, matching that of the constant covariate vector $\mathbf{1}$. The Dantzig selector (Candès and Tao, 2007) is defined as

$$\hat{\beta}_{\text{DS}} = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \|\beta\|_1 \quad \text{subject to} \quad \|n^{-1} \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta)\|_\infty \leq \lambda_1, \quad (2)$$

where $\lambda_1 \geq 0$ is a regularization parameter. The above constant Dantzig selector constraint on correlations between all covariates and the residual vector may not be flexible enough to differentiate important covariates and noise covariates. We introduce an extension of the Dantzig selector, the constrained Dantzig selector, defined as

$$\begin{aligned} \hat{\beta}_{\text{CDS}} = & \operatorname{argmin}_{\beta \in \mathcal{B}_\lambda} \|\beta\|_1 \\ & \text{subject to} \quad |n^{-1} \mathbf{x}_j^T(\mathbf{y} - \mathbf{X}\beta)| \leq \lambda_0 1_{\{|\beta_j| \geq \lambda\}} + \lambda_1 1_{\{|\beta_j| = 0\}} \\ & \text{for } j = 1, \dots, p, \end{aligned} \quad (3)$$

where $\lambda_0 \geq 0$ is a regularization parameter and $\mathcal{B}_\lambda = \{\beta \in \mathbb{R}^p : \beta_j = 0 \text{ or } |\beta_j| \geq \lambda \text{ for each } j\}$ is the constrained parameter space for some $\lambda \geq 0$. When we choose $\lambda = 0$ and $\lambda_0 = \lambda_1$, the constrained Dantzig selector becomes the Dantzig selector. Throughout this paper, we choose the regularization parameters λ_0 and λ_1 as $c_0\{(\log n)/n\}^{1/2}$ and $c_1\{(\log p)/n\}^{1/2}$, respectively, with $\lambda_0 \leq \lambda_1$ as well as c_0 and c_1 two sufficiently large positive constants, and assume that λ is a parameter greater than λ_1 . The two parameters λ_0 and λ_1 differentially bound two types of correlations: on the support of the constrained Dantzig selector, the correlations between covariates and residuals are bounded, up to a common scale, by λ_0 ; on its complement, however, the correlations are bounded through λ_1 . In the ultra-high dimensional case, meaning $\log p = O(n^\alpha)$ for some $0 < \alpha < 1$, the constraints involving λ_0 are tighter than those involving λ_1 , in which λ_1 is a universal regularization parameter for the Dantzig selector; see Candès and Tao (2007) and Bickel et al. (2009).

We now provide more insights into the new constraints in the constrained Dantzig selector. First, it is worthwhile to notice that if $\beta_0 \in \mathcal{B}_\lambda$, β_0 can satisfy new constraints with large probability in model setting (1); see the proof of Theorem 1. With the tighter constraints, the feasible set of the constrained Dantzig selector problem is a subset of that of the Dantzig selector problem, resulting in a search of the solution in a reduced space. Second, it is appealing to extract more information in important covariates, leading to lower correlations between those variables and the residual vector. In this spirit, the constrained Dantzig selector puts tighter constraints on the correlations between selected variables and residuals. Third, the constrained Dantzig selector is defined on the constrained parameter space \mathcal{B}_λ , which has been introduced in Fan and Lv (2013). Such a space also shares some similarity to the union of coordinate subspaces considered in Fan and Lv (2011) for characterizing the restricted global optimality of nonconcave penalized likelihood estimators.

The threshold λ in \mathcal{B}_λ distinguishes between important covariates with strong effects and noise covariates with weak effects. As shown in Fan and Lv (2013), this feature can lead to improved sparsity and effectively prevent overfitting by making it harder for noise covariates to enter the model.

3. Main results

In this section, two parallel theoretical results are introduced based on the uniform uncertainty principle (UUP) condition and restricted eigenvalue assumption, respectively. Although the UUP condition may be relatively stringent in some applications, we still present one theorem for our method under this condition in Section 3.1 for the purpose of comparison with the original Dantzig selector.

3.1 Nonasymptotic compressed sensing properties

Since the Dantzig selector was introduced partly for applications in compressed sensing, we first study the nonasymptotic compressed sensing properties of the constrained Dantzig selector by adopting the theoretical framework in Candès and Tao (2007). They introduced the uniform uncertainty principle condition defined as follows. Denote by \mathbf{X}_T a submatrix of \mathbf{X} consisting of columns with indices in a set $T \subset \{1, \dots, p\}$. For the true model size s , define the s -restricted isometry constant of \mathbf{X} as the smallest constant δ_s such that

$$(1 - \delta_s) \|\mathbf{h}\|_2^2 \leq n^{-1} \|\mathbf{X}_T \mathbf{h}\|_2^2 \leq (1 + \delta_s) \|\mathbf{h}\|_2^2$$

for any set T with size at most s and any vector \mathbf{h} . This condition requires that each submatrix of \mathbf{X} with at most s columns behaves similarly as an orthonormal system. Another constant, the s -restricted orthogonality constant, is defined as the smallest quantity $\theta_{s,2s}$ such that

$$n^{-1} |\langle \mathbf{X}_T \mathbf{h}, \mathbf{X}_{T'} \mathbf{h}' \rangle| \leq \theta_{s,2s} \|\mathbf{h}\|_2 \|\mathbf{h}'\|_2$$

for all pairs of disjoint sets $T, T' \subset \{1, \dots, p\}$ with $|T| \leq s$ and $|T'| \leq 2s$ and any vectors \mathbf{h}, \mathbf{h}' . The uniform uncertainty principle condition is simply stated as

$$\delta_s + \theta_{s,2s} < 1. \quad (4)$$

For notational simplicity, we drop the subscripts and denote these two constants by δ and θ , respectively. Without loss of generality, assume that $\text{supp}(\beta_0) = \{1 \leq j \leq p : \beta_{0,j} \neq 0\} = \{1, \dots, s\}$ hereafter. To evaluate the sparse recovery accuracy, we consider the number of falsely discovered signs defined as $\text{FS}(\hat{\beta}) = |\{j = 1, \dots, p : \text{sgn}(\hat{\beta}_j) \neq \text{sgn}(\beta_{0,j})\}|$ for an estimator $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$. Now we are ready to present the nonasymptotic compressed sensing properties of the constrained Dantzig selector.

Theorem 1 *Assume that the uniform uncertainty principle condition (4) holds and $\beta_0 \in \mathcal{B}_\lambda$ with $\lambda \geq C^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$ for some positive constant C . Then with probability at least*

$1 - O(n^{-c})$ for $c = (c_0^2 \wedge c_1^2)/(2\sigma^2) - 1$, the constrained Dantzig selector $\hat{\beta}$ satisfies that

$$\begin{aligned}\|\hat{\beta} - \beta_0\|_1 &\leq 2\sqrt{5}(1 - \delta - \theta)^{-1}c_0s\sqrt{(\log n)/n}, \\ \|\hat{\beta} - \beta_0\|_2 &\leq \sqrt{5}(1 - \delta - \theta)^{-1}c_0\sqrt{s(\log n)/n}, \\ FS(\hat{\beta}) &\leq Cc_1^2s(\log p)/(n\lambda^2).\end{aligned}$$

If in addition $\lambda > \sqrt{5}(1 - \delta - \theta)^{-1}c_0\sqrt{s(\log n)/n}$, then with the same probability, it also holds that $\text{sgn}(\hat{\beta}) = \text{sgn}(\beta_0)$ and $\|\hat{\beta} - \beta_0\|_\infty \leq 2c_0\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty\sqrt{(\log n)/n}$, where \mathbf{X}_1 is an $n \times s$ submatrix of \mathbf{X} corresponding to s nonzero $\beta_{0,j}$'s.

The constant c in the above probability bound can be sufficiently large since both constants c_0 and c_1 are assumed to be large, while the constant C comes from Theorem 1.1 in Candès and Tao (2007); see the proof in the Appendix for details. In the above bound on the L_∞ -estimation loss, it holds that $\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty \leq s^{1/2}\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_2 \leq (1 - \delta)^{-1}s^{1/2}$. See Section 3.2 for more discussion on this quantity.

From Theorem 1, we see improvements of the constrained Dantzig selector over the Dantzig selector, which has a convergence rate, in terms of the L_2 -estimation loss, up to a factor $\log p$ of that for the ideal procedure. However, the sparsity property of the Dantzig selector was not investigated in Candès and Tao (2007). In contrast, the constrained Dantzig selector is shown to have an inverse quadratic relationship between the number of falsely discovered signs and the threshold λ , revealing that its model selection accuracy increases with the signal strength. The number of falsely discovered signs can be controlled below or as an asymptotically vanishing fraction of the true model size, since $FS(\hat{\beta}) \leq Cs(\lambda_1/\lambda)^2 \leq (1 + \lambda_1/\lambda)^{-2}s < s$ by assuming $\lambda > C^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$.

Another advantage of the constrained Dantzig selector lies in its convergence rates. In the case of ultra-high dimensionality which is typical in compressed sensing applications, its prediction and estimation losses can be reduced from the logarithmic factor $\log p$ to $\log n$ with overwhelming probability. In particular, only a fairly weak assumption on the signal strength is imposed to attain such improved convergence rates. In fact, it has been shown in Raskutti et al. (2011) that without any condition on the signal strength, the minimax convergence rate of L_2 risk has an upper bound of order $O\{s^{1/2}(\log p)/n\}$. Especially, they claimed that the Dantzig selector can achieve such minimax rate, but requires a relatively stronger condition on the design matrix than nonconvex optimization algorithms to determine the minimax upper bounds. We push a step forward that the constrained Dantzig selector, with additional signal strength conditions, can attain the L_2 -estimation loss of a smaller order than $O\{s^{1/2}(\log p)/n\}$ in ultra-high dimensions.

There exist other methods which have been shown to enjoy convergence rates of the same order as well, for example, in Zheng et al. (2014) for high-dimensional thresholded regression. However, these results usually rely on a stronger condition on signal strength, such as, the minimum signal strength is at least of order $\{s(\log p)/n\}^{1/2}$. In another work, Fan and Lv (2011) showed that the nonconcave penalized estimator can have a consistency rate of $O_p(s^{1/2}n^{-\gamma}\log n)$ for some $\gamma \in (0, 1/2]$ under the L_2 -estimation loss, which can be slower than our rate of convergence. More detailed discussion on the relationship between the faster rates of convergence and the assumptions on the signal strength and sparsity can be found in Section 3.2. A main implication of our improved convergence rates is that

a smaller number of observations will be needed for the constrained Dantzig selector to attain the same level of accuracy in compressed sensing, as the Dantzig selector, which is demonstrated in Section 4.2.

3.2 Sampling properties

The properties of the Dantzig selector have also been extensively investigated in Bickel et al. (2009). They introduced the restricted eigenvalue assumption with which the oracle inequalities under various prediction and estimation losses were derived. We adopt their theoretical framework and study the sampling properties of the constrained Dantzig selector under the restricted eigenvalue assumption stated below. A positive integer m is said to be in the same order of s if m/s can be bounded from both above and below by some positive constants.

Condition 1 *For some positive integer m in the same order of s , there exists some positive constant κ such that $\|n^{-1/2}\mathbf{X}\boldsymbol{\delta}\|_2 \geq \kappa \max\{\|\boldsymbol{\delta}_1\|_2, \|\boldsymbol{\delta}'_1\|_2\}$ for all $\boldsymbol{\delta} \in \mathbb{R}^p$ satisfying $\|\boldsymbol{\delta}_2\|_1 \leq \|\boldsymbol{\delta}_1\|_1$, where $\boldsymbol{\delta} = (\boldsymbol{\delta}_1^T, \boldsymbol{\delta}_2^T)^T$, $\boldsymbol{\delta}_1$ is a subvector of $\boldsymbol{\delta}$ consisting of the first s components, and $\boldsymbol{\delta}'_1$ is a subvector of $\boldsymbol{\delta}_2$ consisting of the $\max\{m, C_m s(\lambda_1/\lambda)^2\}$ largest components in magnitude, with C_m some positive constant.*

Condition 1 is a basic assumption on the design matrix \mathbf{X} for deriving the oracle inequalities of the Dantzig selector. Since we assume that $\text{supp}(\boldsymbol{\beta}_0) = \{1, \dots, s\}$, Condition 1 indeed plays the same role as the restricted eigenvalue assumption $\text{RE}(s, m, 1)$ in Bickel et al. (2009), which assumes the inequality in Condition 1 holds for any subset with size no larger than s to cover all possibilities of $\text{supp}(\boldsymbol{\beta}_0)$. See (10) in the Appendix for insights into the basic inequality $\|\boldsymbol{\delta}_2\|_1 \leq \|\boldsymbol{\delta}_1\|_1$ and Bickel et al. (2009) for more detailed discussions on this assumption.

Theorem 2 *Assume that Condition 1 holds and $\boldsymbol{\beta}_0 \in \mathcal{B}_\lambda$ with $\lambda \geq C_m^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$. Then the constrained Dantzig selector $\hat{\boldsymbol{\beta}}$ satisfies with the same probability as in Theorem 1 that*

$$\begin{aligned} n^{-1/2}\|\mathbf{X}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)\|_2 &= O(\kappa^{-1}\sqrt{s(\log n)/n}), & \|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_1 &= O(\kappa^{-2}s\sqrt{(\log n)/n}), \\ \|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_2 &= O(\kappa^{-2}\sqrt{s(\log n)/n}), & FS(\hat{\boldsymbol{\beta}}) &\leq C_m c_1^2 s(\log p)/(n\lambda^2). \end{aligned}$$

If in addition $\lambda > 2\sqrt{5}\kappa^{-2}c_0s^{1/2}\sqrt{(\log n)/n}$, then with the same probability it also holds that $\text{sgn}(\hat{\boldsymbol{\beta}}) = \text{sgn}(\boldsymbol{\beta}_0)$ and $\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0\|_\infty = O\{\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty\sqrt{(\log n)/n}\}$.

Theorem 2 establishes asymptotic results on the sparsity and oracle inequalities for the constrained Dantzig selector under the restricted eigenvalue assumption. This assumption, which is an alternative to the uniform uncertainty principle condition, has also been widely employed in high-dimensional settings. In Bickel et al. (2009), an approximate equivalence of the Lasso estimator (Tibshirani, 1996) and Dantzig selector was proved under this assumption, and the Lasso estimator was shown to be sparse with size $O(\phi_{\max}s)$, where ϕ_{\max} is the largest eigenvalue of the Gram matrix $n^{-1}\mathbf{X}^T\mathbf{X}$. In contrast, the constrained Dantzig selector gives a sparser model under the restricted eigenvalue assumption, since its

number of falsely discovered signs $\text{FS}(\hat{\beta}) \leq C_m s (\lambda_1/\lambda)^2 = o(s)$ when $\lambda_1 = o(\lambda)$. Similar as in Theorem 1, the constrained Dantzig selector improves over both Lasso and the Dantzig selector in terms of convergence rates, a reduction of the $\log p$ factor to $\log n$.

Let us now take a closer look at the relationship between the faster rates of convergence and the assumptions on the signal strength and sparsity. Observe that the enhanced consistency results require the minimum signal strength to be at least of order $(\log p)/\sqrt{n \log n}$, in view of the assumption $\lambda \geq C_m^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$. Assume for simplicity that $\|\beta_0\|_2$ is bounded away from both zero and ∞ . Then the order of the minimum signal strength yields an upper bound on the sparsity level $s = O\{n(\log n)/(\log p)^2\}$, which means that the sparsity s is required to be smaller when the dimensionality p becomes larger. In the ultra-high dimensional setting of $\log p = O(n^\alpha)$ with $0 < \alpha < 1$, we then have $s = O(n^{1-2\alpha} \log n)$. Thus the classical convergence rates involving $s(\log p)/n = O(n^{-\alpha} \log n)$ still go to zero asymptotically, while the rates established in our paper are improved with $s(\log p)/n$ replaced by $s(\log n)/n = O\{n^{-2\alpha}(\log n)^2\}$. We see a gain on the convergence rate of a factor $n^\alpha/(\log n)$ at the cost of the aforementioned signal strength and sparsity assumptions.

One can see that results in Theorems 1 and 2 are approximately equivalent, while the latter presents an additional oracle inequality on the prediction loss. An interesting phenomenon is that if adopting a simpler version, that is, the Dantzig selector equipped with the thresholding constraint only, we can also obtain similar results as in Theorems 1 and 2, but with a stronger condition on signal strength such as $\lambda \gg s^{1/2}\lambda_1$. In this sense, the constrained Dantzig selector is also an extension of the Dantzig selector equipped with the thresholding constraint only, but enjoys better properties. Some comprehensive results on the prediction and variable selection properties have also been established in Fan and Lv (2013) for various regularization methods, revealing their asymptotic equivalence in the thresholded parameter space. However, as mentioned in Section 3.1, improved rates as in Theorem 2 commonly require a stronger assumption on the signal strength, which is $\beta_0 \in \mathcal{B}_\lambda$ with $\lambda \gg s^{1/2}\lambda_1$; see, for example, Theorem 2 of Fan and Lv (2013).

For the quantity $\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty$ in the above bound on the L_∞ -estimation loss, if \mathbf{X}_1 takes the form of a common correlation matrix $(1 - \rho)I_s + \rho\mathbf{1}_s\mathbf{1}_s^T$ for some $\rho \in [0, 1)$, it is easy to check that $\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty = (1 - \rho)^{-1}\{1 + (2s - 3)\rho\}/\{1 + (s - 1)\rho\}$, which is bounded regardless of the value of s .

3.3 Asymptotic properties of computable solutions

The nonasymptotic and sampling properties of the constrained Dantzig selector, as the global minimizer, have been established in Sections 3.1 and 3.2, respectively. However, it is not guaranteed that the global minimizer can be generated by a computational algorithm. Moreover, a computable solution, generated by any algorithm, may only be a local minimizer in many cases. Under certain regularity conditions, we demonstrate that the local minimizer of our method can still share the same nice asymptotic properties as the global one.

Theorem 3 *Let $\hat{\beta}$ be a computable local minimizer of (3). Assume that there exist some positive constants c_2, κ_0 and sufficiently large positive constant c_3 such that $\|\hat{\beta}\|_0 \leq c_2 s$ and $\min_{\|\delta\|_2=1, \|\delta\|_0 \leq c_3 s} n^{-1/2}\|\mathbf{X}\delta\|_2 \geq \kappa_0$. Then under conditions of Theorem 2, $\hat{\beta}$ enjoys the same properties as the global minimizer in Theorem 2.*

In Section 4.1, we introduce an efficient algorithm that gives us a local minimizer. Theorem 3 indicates that the obtained solution can also enjoy the asymptotic properties in Theorem 2 under the extra assumptions that $\hat{\beta}$ is a sparse solution with the number of nonzero components comparable with s and the design matrix \mathbf{X} satisfies a sparse eigenvalue condition. Similar results for the computable solution can be found in Fan and Lv (2013, 2014), where the local minimizer is additionally assumed to satisfy certain constraint on the correlation between the residual vector and all the covariates.

4. Numerical studies

In this section, we first introduce an algorithm which can efficiently implement the constrained Dantzig selector. Then several simulation studies and two real data examples are presented to evaluate the performance of our method.

4.1 Implementation

The constrained Dantzig selector defined in (3) depends on tuning parameters λ_0 , λ_1 , and λ . We suggest some fixed values for λ_0 and λ to simplify the computation, since the proposed method is generally not that sensitive to λ_0 and λ as long as they fall in certain ranges. In simulation studies to be presented, a value around $\{(\log p)/n\}^{1/2}$ for λ and a smaller value for λ_0 , say $0.05\{(\log p)/n\}^{1/2}$ or $0.1\{(\log p)/n\}^{1/2}$, can provide us nice prediction and estimation results. The value of λ_0 is chosen to be smaller than $\{(\log p)/n\}^{1/2}$ to mitigate the selection of noise variables and facilitate sparse modeling. The performance of our method with respect to different values of λ_0 and λ is shown in simulation Example 2 in Section 4.2, which is a typical example that illustrates the robustness of the proposed method with respect to λ_0 and λ .

For fixed λ_0 and λ , we exploit the idea of sequential linear programming to produce the solution path of the constrained Dantzig selector as λ_1 varies. Choose a grid of values for the tuning parameter λ_1 in decreasing order with the first one being $\|n^{-1}\mathbf{X}^T\mathbf{y}\|_\infty$. It is easy to check that $\beta = \mathbf{0}$ satisfies all the constraints in (3) for $\lambda_1 = \|n^{-1}\mathbf{X}^T\mathbf{y}\|_\infty$, and thus the solution is $\hat{\beta}_{\text{CDS}} = \mathbf{0}$ in this case. For each λ_1 in the grid, we use the solution from the previous one in the grid as an initial value to speed up the convergence. For a given λ_1 , we define an active set, iteratively update this set, and solve the constrained Dantzig selector problem. We name this algorithm as the CDS algorithm which is detailed in four steps below.

1. For a fixed λ_1 in the grid, denote by $\hat{\beta}_{\lambda_1}^{(0)}$ the initial value. Let $\hat{\beta}_{\lambda_1}^{(0)}$ be zero when $\lambda_1 = \|n^{-1}\mathbf{X}^T\mathbf{y}\|_\infty$, and the estimate from previous λ_1 in the grid otherwise.
2. Denote by $\hat{\beta}_{\lambda_1}^{(k)}$ the estimate from the k th iteration. Define the active set \mathcal{A} as the support of $\hat{\beta}_{\lambda_1}^{(k)}$ and \mathcal{A}^c its complement. Let \mathbf{b} be a vector with constant components λ_0 on \mathcal{A} and λ_1 on \mathcal{A}^c . For the $(k+1)$ th iteration, update \mathcal{A} as $\mathcal{A} \cup \{j \in \mathcal{A}^c : |n^{-1}\mathbf{x}_j^T(\mathbf{y} - \mathbf{X}\hat{\beta}_{\mathcal{A}})| > \lambda_1\}$, where the subscript \mathcal{A} indicates a subvector restricted on

\mathcal{A} . Solve the following linear program on the new set \mathcal{A} :

$$\hat{\beta}_{\mathcal{A}} = \operatorname{argmin} \|\beta_{\mathcal{A}}\|_1 \text{ subject to } |n^{-1}\mathbf{X}_{\mathcal{A}}^T(\mathbf{y} - \mathbf{X}_{\mathcal{A}}\beta_{\mathcal{A}})| \preceq \mathbf{b}_{\mathcal{A}}, \quad (5)$$

where \preceq is understood as componentwise no larger than and the subscript \mathcal{A} also indicates a submatrix with columns corresponding to \mathcal{A} . For the solution obtained in (5), set all its components smaller than λ in magnitude to zero.

3. Update the active set \mathcal{A} as the support of $\hat{\beta}_{\mathcal{A}}$. Solve the Dantzig selector problem on this active set with λ_0 as the regularization parameter:

$$\hat{\beta}_{\mathcal{A}} = \operatorname{argmin} \|\beta_{\mathcal{A}}\|_1 \text{ subject to } \|n^{-1}\mathbf{X}_{\mathcal{A}}^T(\mathbf{y} - \mathbf{X}_{\mathcal{A}}\beta_{\mathcal{A}})\|_{\infty} \leq \lambda_0. \quad (6)$$

Let $\hat{\beta}_{\mathcal{A}}^{(k+1)} = \hat{\beta}_{\mathcal{A}}$ and $\hat{\beta}_{\mathcal{A}^c}^{(k+1)} = \mathbf{0}$, which give the solution for the $(k+1)$ th iteration.

4. Repeat steps 2 and 3 until convergence for a fixed λ_1 and record the estimate from the last iteration as $\hat{\beta}_{\lambda_1}$. Jump to the next λ_1 if $\hat{\beta}_{\lambda_1} \in \mathcal{B}_{\lambda}$, and stop the algorithm otherwise.

With the solution path produced, we use the cross-validation to select the tuning parameter λ_1 . One can also tune λ_0 and λ similarly as for λ_1 , but as suggested before, some fixed values for them suffice to obtain satisfactory results.

The rationales of the constrained Dantzig selector algorithm are as follows. Step 1 defines the initial value (0th iteration) for each λ_1 in the grid. In step 2, starting with a smaller active set, we add variables that violate the constrained Dantzig selector constraints to eliminate such conflict. As a consequence, some components of $\mathbf{b}_{\mathcal{A}}$ are of value λ_1 instead of λ_0 . Therefore, we need to further solve (6) in step 3 by noting that restricted on its support, the constrained Dantzig selector should be a solution to the Dantzig selector problem with parameter λ_0 . An early stopping of the solution path is imposed in step 4 to make this algorithm computationally more efficient.

An appealing feature of this algorithm is that its convergence can be checked easily. Once there are no more variables violating the constrained Dantzig selector constraints, that is, $\{j \in \mathcal{A}^c : |n^{-1}\mathbf{x}_j^T(\mathbf{y} - \mathbf{X}\hat{\beta}_{\mathcal{A}})| > \lambda_1\} = \emptyset$, the iteration stops and the algorithm converges. In other words, the convergence of the algorithm is equivalent to that of the active set which can be checked directly. When the algorithm converges, the solution lies in the feasible set of the constrained Dantzig selector problem and is a global minimizer restricted on the active set, and is thus a local minimizer.

In simulation Example 2 of Section 4.2, we tracked the convergence property of the algorithm on 100 data sets for $p = 1000, 5000$, and 10000 , respectively. In all cases, we observe that the algorithm always converged over all 100 simulations, indicating considerable stability of this algorithm. Another advantage of the algorithm is that it is built upon the Dantzig selector in lower dimensions, so it inherits the computational efficiency.

4.2 Simulation studies

To better illustrate the performance of the constrained Dantzig selector, we consider the thresholded Dantzig selector which simply sets components of the Dantzig selector estimate to zeros if smaller than a threshold in magnitude. We evaluated the performance

of the constrained Dantzig selector in comparison with the Dantzig selector, thresholded Dantzig selector, Lasso, elastic net (Zou and Hastie, 2005), and adaptive Lasso (Zou, 2006). Two simulation studies were considered, with the first one investigating sparse recovery for compressed sensing and the second one examining sparse modeling.

The setting of the first simulation example is similar to that of the sparse recovery example in Lv and Fan (2009). The noiseless sparse recovery example is considered here since the Dantzig selector problem originated from compressed sensing. We want to evaluate the capability of our constrained Dantzig selector in recovering sparse signals as well. We generated 100 data sets from model (1) without noise, that is, the linear equation $\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_0$ with $(s, p) = (7, 1000)$. The nonzero components of $\boldsymbol{\beta}_0$ were set to be $(1, -0.5, 0.7, -1.2, -0.9, 0.3, 0.55)^T$ lying in the first seven components, and n was chosen to be even integers between 30 and 80. The rows of the design matrix \mathbf{X} were sampled as independent and identically distributed (i.i.d.) copies from $N(\mathbf{0}, \boldsymbol{\Gamma}_r)$, with $\boldsymbol{\Gamma}_r$ a $p \times p$ matrix with diagonal elements being 1 and off-diagonal elements being r , and then each column was rescaled to have L_2 -norm \sqrt{n} . Three levels of population collinearity, $r = 0, 0.2$, and 0.5 , were considered. Let λ_0 and λ be in two small grids $\{0.001, 0.005, 0.01, 0.05, 0.1\}$ and $\{0.05, 0.1, 0.15, 0.2\}$, respectively. We chose two grids of values for them since in the literature of compressed sensing, it is desirable to have the true support included among a set of estimates. The value 0.2 was chosen because it is close to $\{(\log p)/n\}^{1/2}$, which is about $\{(\log 1000)/80\}^{1/2}$ in this example. Smaller values for λ and λ_0 are also included in the grid for conservativeness. We set the grid of values for λ_1 as described in Section 4.1. If any of the solutions in the path had exactly the same support as $\boldsymbol{\beta}_0$, it is counted as successful recovery. This criterion was applied to all other methods in this example for fair comparison.

Figure 1 presents the probabilities of exact recovery of sparse $\boldsymbol{\beta}_0$ based on 100 simulations by all methods. We see that all methods performed well in relatively large samples and had lower probability of successful sparse recovery when the sample size becomes smaller. The constrained Dantzig selector performed better than other methods over different sample sizes and three levels of population collinearity. In particular, the thresholded Dantzig selector performed similarly to the Dantzig selector, revealing that simple thresholding alone, instead of flexible constraints as in the constrained Dantzig selector, does not help much on signal recovery in this case.

The second simulation example adopts a similar setting to that in Zheng et al. (2014). We generated 100 data sets from the linear regression model (1) with Gaussian error $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 I_n)$. The coefficient vector $\boldsymbol{\beta}_0 = (\mathbf{v}^T, \dots, \mathbf{v}^T, \mathbf{0}^T)^T$ with the pattern $\mathbf{v} = (\boldsymbol{\beta}_{\text{strong}}^T, \boldsymbol{\beta}_{\text{weak}}^T)^T$ repeated three times, where $\boldsymbol{\beta}_{\text{strong}} = (0.6, 0, 0, -0.6, 0, 0)^T$ and $\boldsymbol{\beta}_{\text{weak}} = (0.05, 0, 0, -0.05, 0, 0)^T$. The coefficient subvectors $\boldsymbol{\beta}_{\text{strong}}$ and $\boldsymbol{\beta}_{\text{weak}}$ stand for the strong signals and weak signals in $\boldsymbol{\beta}_0$, respectively. The sample size and noise level were chosen as $(n, \sigma) = (100, 0.4)$, while the dimensionality p was set to be 1000, 5000, and 10000. The rows of the $n \times p$ design matrix \mathbf{X} were sampled as i.i.d. copies from a multivariate normal distribution $N(\mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = (0.5^{|i-j|})_{1 \leq i, j \leq p}$. We applied all methods as in simulation example 1 and set $\lambda_0 = 0.01$ and $\lambda = 0.2$ for our method. Similarly as before, the value of 0.2 was selected since it is close to $\{(\log 1000)/100\}^{1/2}$. The ideal procedure, which knows the true underlying sparse model in advance, was also used as a benchmark.




Figure 1: Probabilities of exact sparse recovery for the Dantzig selector (DS), thresholded Dantzig selector (TDS), Lasso, elastic net (Enet), adaptive Lasso (ALasso), and constrained Dantzig selector (CDS) in simulation example 1 of Section 4.2.

To compare these methods, we considered several performance measures: the prediction error, the L_q -estimation loss with $q = 1, 2, \infty$, number of false positives, and number of false negatives for strong or weak signals. The prediction error is defined as $E(Y - \mathbf{x}^T \hat{\beta})^2$ with $\hat{\beta}$ an estimate and (\mathbf{x}^T, Y) an independent observation, and the expectation was calculated using an independent test sample of size 10000. A false positive means a falsely selected noise covariate in the model, while a false negative means a missed true covariate. Table 1 summarizes the comparison results by all methods. We observe that most weak covariates were missed by all methods. This is reasonable since the weak signals are around the noise level, making it difficult to distinguish them from the noise covariates. However, the constrained Dantzig selector outperformed other methods in terms of other prediction and estimation measures, and followed very closely the ideal procedure in all cases of $p = 1000, 5000$, and 10000 . In particular, the L_∞ -estimation loss for the constrained Dantzig selector was similar to that for the oracle procedure, confirming tight bounds on this loss in Theorems 1 and 2. As the dimensionality grows higher, the constrained Dantzig selector performed similarly, while other methods suffered from high dimensionality. In particular, the thresholded Dantzig selector has been shown to improve over the Dantzig selector, but was still outperformed by the adaptive Lasso and constrained Dantzig selector in this example, revealing the necessity to introduce more flexible constraints instead of simple thresholding.

Recall that we have fixed $\lambda_0 = 0.01$ and $\lambda = 0.2$ in the simulation example 2 across all settings. We now study the robustness of the constrained Dantzig with respect to λ_0 and λ in this typical example. For simplicity, we only consider Example 2 with dimensionality $p = 1000$. Instead of fixing $\lambda_0 = 0.01$ and $\lambda = 0.2$, we let λ_0 be a value in the grid $\{0.001, 0.005, 0.01, 0.015, 0.02, 0.025, 0.03\}$ and λ a value in $\{0.1, 0.15, 0.2, 0.25, 0.3, 0.35\}$, respectively. Therefore, we study $7 \times 6 = 42$ different combinations of λ_0 and λ in total to

Table 1: Means and standard errors (in parentheses) of different performance measures by all methods in simulation example 2 of Section 4.2.

Measure	DS	TDS	Lasso	Enet	ALasso	CDS	Oracle
$p = 1000$							
PE ($\times 10^{-2}$)	30.8 (0.6)	28.5 (0.5)	30.3 (0.5)	32.9 (0.7)	19.1 (0.1)	18.5 (0.1)	18.2 (0.1)
L_1 ($\times 10^{-2}$)	201.9 (5.4)	137.2 (3.4)	186.1 (5.0)	211.1 (6.0)	58.0 (1.1)	51.3 (0.6)	41.5 (0.9)
L_2 ($\times 10^{-2}$)	40.1 (0.7)	37.2 (0.7)	39.8 (0.7)	43.0 (0.8)	18.3 (0.3)	16.3 (0.2)	14.7 (0.3)
L_∞ ($\times 10^{-2}$)	19.0 (0.5)	18.6 (0.4)	19.3 (0.5)	20.7 (0.5)	9.1 (0.3)	7.5 (0.2)	8.4 (0.2)
FP	44.4 (1.8)	5.5 (3.8)	36.3 (1.6)	44.1 (1.9)	0.5 (0.1)	0 (0)	0 (0)
FN.strong	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
FN.weak	5.3 (0.1)	5.9 (0.4)	5.4 (0.1)	5.3 (0.1)	6.0 (0.0)	6.0 (0.0)	0 (0)
$p = 5000$							
PE ($\times 10^{-2}$)	45.1 (1.1)	39.3 (1.1)	44.8 (1.1)	44.9 (1.1)	21.3 (0.6)	18.4 (0.1)	18.3 (0.1)
L_1 ($\times 10^{-2}$)	289.3 (6.4)	184.6 (4.5)	270.8 (6.8)	273.2 (6.9)	71.2 (2.1)	50.4 (0.7)	41.7 (1.1)
L_2 ($\times 10^{-2}$)	56.3 (1.1)	50.6 (1.1)	56.1 (1.1)	56.2 (1.1)	22.9 (0.9)	16.0 (0.2)	14.9 (0.4)
L_∞ ($\times 10^{-2}$)	27.4 (0.7)	25.1 (0.6)	27.8 (0.7)	27.8 (0.7)	12.7 (0.7)	7.3 (0.2)	8.8 (0.3)
FP	60.6 (1.7)	7.0 (3.5)	53.5 (2.3)	53.8 (2.1)	1.0 (0.1)	0 (0)	0 (0)
FN.strong	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)
FN.weak	5.2 (0.1)	6.0 (0.1)	5.4 (0.1)	5.3 (0.1)	5.9 (0.0)	6.0 (0.0)	0 (0)
$p = 10000$							
PE ($\times 10^{-2}$)	54.4 (16.2)	54.8 (17.4)	56.7 (23.1)	63.3 (21.7)	32.4 (32.3)	19.0 (0.5)	18.3 (0.1)
L_1 ($\times 10^{-2}$)	322.3 (67.9)	218 (60.3)	296.3 (72.7)	334.8 (79.9)	86.6 (67.6)	51.9 (1.6)	41.3 (0.9)
L_2 ($\times 10^{-2}$)	62.7 (14.1)	63.5 (14.9)	64.2 (17.7)	70.3 (16.8)	29.1 (26.3)	16.6 (0.6)	14.7 (0.3)
L_∞ ($\times 10^{-2}$)	30.5 (8.2)	32.5 (8.5)	31.8 (9.5)	34.4 (9.3)	15.8 (13.2)	7.8 (0.6)	8.5 (0.2)
FP	69.9 (14.3)	6.9 (4.6)	56.89 (20.66)	64.06 (18.64)	0.78 (1.25)	0.01 (0.10)	0 (0)
FN.strong	0 (0.1)	0 (0.1)	0.02 (0.14)	0.02 (0.14)	0.21 (0.95)	0.01 (0.10)	0 (0)
FN.weak	6 (0.2)	6 (0.1)	5.98 (0.14)	5.97 (0.17)	6 (0)	6 (0)	0 (0)

evaluate the robustness. The same performance measures were calculated. Here we only present the prediction error results to save space, and the other results are available upon request. The tuning parameter λ_1 was chosen by cross-validation, similar as in Example 2. We can see from Table 2 that for all $\lambda \geq 0.15$ and all λ_0 in the grids, the means and standard errors of the prediction error are very close or even identical to those for $\lambda_0 = 0.01$ and $\lambda = 0.2$. The prediction error in the case of $\lambda = 0.1$ is slightly higher since when the threshold becomes lower, some noise variables can be included. For the results on estimation losses as well as false positives and false negatives, we observe similar patterns confirming the robustness of our method with respect to the choices of λ_0 and λ .

4.3 Real data analyses

We applied the same methods as in Section 4.2 to two real data sets: one real PCR data set and another gene expression data set, both in the high-dimensional setting with relatively small sample size. In both data sets, we found that the proposed method enjoys smaller prediction errors and the differences are statistically significant.

Table 2: Means and standard errors (in parentheses) of prediction error of the constrained Dantzig selector for different choices of λ_0 and λ in simulation example 2 with $p = 1000$.

λ_0	λ					
	0.10	0.15	0.20	0.25	0.30	0.35
0.001	0.192 (0.002)	0.186 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)
0.005	0.191 (0.003)	0.188 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)
0.010	0.192 (0.002)	0.187 (0.001)	0.185 (0.001)	0.185 (0.001)	0.188 (0.003)	0.185 (0.001)
0.015	0.196 (0.007)	0.186 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)
0.020	0.192 (0.002)	0.188 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)
0.025	0.190 (0.002)	0.187 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)	0.185 (0.001)
0.030	0.191 (0.002)	0.187 (0.001)	0.191 (0.007)	0.188 (0.003)	0.185 (0.001)	0.185 (0.001)

The real PCR data set, originally studied in Lan et al. (2006), examines the genetics of two inbred mouse populations. This data set is comprised of $n = 60$ samples with 29 males and 31 females. Expression levels of 22,575 genes were measured. Following Song and Liang (2015), we study the linear relationship between the numbers of Phosphoenolpyruvate carboxykinase (PEPCK), a phenotype measured by quantitative real-time PCR, and the gene expression levels. Both the phenotype and predictors are continuous. As suggested in Song and Liang (2015), we only picked $p = 2000$ genes having the highest marginal correlations with PEPCK as predictors. The response was standardized to have zero mean and unit variance before we conducted the analysis. The 2000 predictors were standardized to have zero mean and L_2 -norm \sqrt{n} in each column. Then the data set was randomly split into a training set of 55 samples and a test set with the remaining 5 samples for 100 times. We set $\lambda_0 = 0.001$ and $\lambda = 0.02$ in this real data analysis as well as the other one below for conservativeness. Methods under comparison are the same as in the simulation studies.

Table 3 reports the means and standard errors of the prediction error on the test data as well as the median model size for each method. We see that the method CDS gave the lowest mean prediction error. Paired t -tests were conducted for the prediction errors of CDS versus DS, TDS, Lasso, ALasso, and Enet, respectively, to test the differences in performance across various methods. The corresponding p -values were 0.0243, 0.0014, 0.0081, 0.0001, and 0.0102, respectively, indicating significantly different prediction error from that for CDS.

The second data set has been studied in Scheetz et al. (2006) and Huang et al. (2008). In this data set, 120 twelve-week-old male rats were selected for tissue harvesting from the eyes and for microarray analysis. There are 31,042 different probe sets in the microarrays from the RNA of those rat eyes. Following Huang et al. (2008), we excluded the probes that were not expressed sufficiently or that lacked sufficient variation, leaving 18,976 probes which satisfy these two criteria. The response variable TRIM32, which was recently found to cause Bardet-Biedl syndrome, is one of the selected 18,976 probes. We then selected 3,000 probes with the largest variances from the remaining 18,975 probes. The goal of our analysis is to identify the genes that are most relevant to the expression level of TRIM32 from the 3,000 candidate genes.

Table 3: Means and standard errors of the prediction error and median model size over 100 random splits for the real PCR data set.

Method	PE	Model Size
DS	0.773 (0.040)	54
TDS	0.897 (0.063)	14
Lasso	0.802 (0.043)	58
ALasso	0.922 (0.056)	5
Enet	0.793 (0.041)	58
CDS	0.660 (0.036)	21

Similarly as in the analysis of the first real data set, we standardized the response and predictors beforehand. The training set contains 100 samples and was sampled randomly 100 times from the full data set. The remaining 20 samples at each time served as the test set. Results of prediction errors and median model sizes are presented in Table 4. It is clear that the proposed method CDS enjoys the lowest prediction error with a small model size. We conducted the same paired t -tests as in the real PCR data set for comparison. The corresponding p -values were 0.0351, 0.0058, 0.0164, 0.0004, and 0.0344, respectively, showing significant improvement.

Table 4: Means and standard errors of the prediction error and median model size over 100 random splits for the gene expression data set of rat eyes.

Method	PE	Model Size
DS	0.582 (0.044)	28.5
TDS	0.627 (0.051)	9
Lasso	0.590 (0.043)	33
ALasso	0.652 (0.051)	8.5
Enet	0.576 (0.040)	67.5
CDS	0.520 (0.025)	9

5. Discussion

We have shown that the suggested constrained Dantzig selector can achieve convergence rates within a logarithmic factor of the sample size of the oracle rates in ultra-high dimensions under a fairly weak assumption on the signal strength. Our work provides a partial answer to an interesting question of whether convergence rates involving a logarithmic factor of the dimensionality are optimal for regularization methods in ultra-high dimensions. It would be interesting to investigate such a phenomenon for more general regularization methods.

Our formulation of the constrained Dantzig selector uses the L_1 -norm of the parameter vector. A natural extension of the method is to exploit the weighted L_1 -norm of the

parameter to allow for different regularization on different covariates, as is in the adaptive Lasso (Zou, 2006). It would be interesting to investigate the behavior of these methods in more general model settings including generalized linear models and survival analysis. These problems are beyond the scope of the current paper and will be interesting topics for future research.

Appendix: Proofs of main results

Proof of Theorem 1

High probability event \mathcal{E} . Recall that $\lambda_0 = c_0 \sqrt{(\log n)/n}$ and $\lambda_1 = c_1 \sqrt{(\log p)/n}$. All the results in Theorems 1 and 2 will be shown to hold on a key event

$$\mathcal{E} = \{\|n^{-1}\mathbf{X}_1^T \boldsymbol{\varepsilon}\|_\infty \leq \lambda_0 \text{ and } \|n^{-1}\mathbf{X}_2^T \boldsymbol{\varepsilon}\|_\infty \leq \lambda_1\}, \quad (7)$$

where \mathbf{X}_1 is a submatrix of \mathbf{X} consisting of columns corresponding to $\text{supp}(\boldsymbol{\beta}_0)$ and \mathbf{X}_2 consists of the remaining columns. Thus we will have the same probability bound in both theorems. The probability bound on the event \mathcal{E} in (7) can be easily calculated, using the classical Gaussian tail probability bound (see, for example, Dudley, 1999) and the Bonferroni inequality, as

$$\begin{aligned} \text{pr}(\mathcal{E}) &\geq 1 - \{\text{pr}(\|n^{-1}\mathbf{X}_1^T \boldsymbol{\varepsilon}\|_\infty > \lambda_0) + \text{pr}(\|n^{-1}\mathbf{X}_2^T \boldsymbol{\varepsilon}\|_\infty > \lambda_1)\} \\ &= 1 - \{s(2/\pi)^{1/2} \sigma \lambda_0^{-1} n^{-1/2} e^{-\lambda_0^2 n / (2\sigma^2)} + (p-s)(2/\pi)^{1/2} \sigma \lambda_1^{-1} n^{-1/2} e^{-\lambda_1^2 n / (2\sigma^2)}\} \\ &= 1 - O\{s n^{-c_0^2/(2\sigma^2)} (\log n)^{-1/2} + (p-s) p^{-c_1^2/(2\sigma^2)} (\log p)^{-1/2}\}, \end{aligned} \quad (8)$$

where the last equality follows from the definitions of λ_0 and λ_1 . Let $c = (c_0^2 \wedge c_1^2)/(2\sigma^2) - 1$ be a sufficiently large positive constant, since the two positive constants c_0 and c_1 are chosen large enough. Recall that p is understood implicitly as $\max\{n, p\}$ throughout the paper. Thus it follows from (8), $s \leq n$, and $n \leq p$ that

$$\text{pr}(\mathcal{E}) = 1 - O(n^{-c}). \quad (9)$$

From now on, we derive all the bounds on the event \mathcal{E} . In particular, in light of (7) and $\boldsymbol{\beta}_0 \in \mathcal{B}_\lambda$ it is easy to verify that conditional on \mathcal{E} , the true regression coefficient vector $\boldsymbol{\beta}_0$ satisfies the constrained Dantzig selector constraints; in other words, $\boldsymbol{\beta}_0$ lies in the feasible set in (3).

Nonasymptotic properties of $\hat{\boldsymbol{\beta}}$. We first make a simple observation on the constrained Dantzig selector $\hat{\boldsymbol{\beta}} = (\hat{\beta}_1, \dots, \hat{\beta}_p)^T$. Recall that without loss of generality, we assume $\text{supp}(\boldsymbol{\beta}_0) = \{1, \dots, s\}$. Let $\boldsymbol{\beta}_0 = (\boldsymbol{\beta}_1^T, \mathbf{0}^T)^T$ with each component of $\boldsymbol{\beta}_1$ being nonzero, and $\hat{\boldsymbol{\beta}} = (\hat{\boldsymbol{\beta}}_1^T, \hat{\boldsymbol{\beta}}_2^T)^T$ with $\hat{\boldsymbol{\beta}}_1$ a subvector of $\hat{\boldsymbol{\beta}}$ consisting of its first s components. Denote by $\boldsymbol{\delta} = (\boldsymbol{\delta}_1^T, \boldsymbol{\delta}_2^T)^T = \hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0$ the estimation error, where $\boldsymbol{\delta}_1 = \hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1$ and $\boldsymbol{\delta}_2 = \hat{\boldsymbol{\beta}}_2$. It follows from the global optimality of $\hat{\boldsymbol{\beta}}$ that $\|\hat{\boldsymbol{\beta}}_1\|_1 + \|\hat{\boldsymbol{\beta}}_2\|_1 = \|\hat{\boldsymbol{\beta}}\|_1 \leq \|\boldsymbol{\beta}_0\|_1 = \|\boldsymbol{\beta}_1\|_1$, which entails that

$$\|\boldsymbol{\delta}_2\|_1 = \|\hat{\boldsymbol{\beta}}_2\|_1 \leq \|\boldsymbol{\beta}_1\|_1 - \|\hat{\boldsymbol{\beta}}_1\|_1 \leq \|\hat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1\|_1 = \|\boldsymbol{\delta}_1\|_1. \quad (10)$$

We will see that this basic inequality $\|\delta_2\|_1 \leq \|\delta_1\|_1$ plays a key role in the technical derivations of both Theorem 1 and 2. Equipped with this inequality and conditional on \mathcal{E} , we are now able to start the derivation of all results in Theorem 1 as follows.

The main idea is to first prove its sparsity property which will be presented in the next paragraph. Then, with the control on the number of false positives and false negatives, we derive an upper bound for the L_2 -estimation loss using the conclusion in Lemma 3.1 of Candès and Tao (2007). Results on other types of losses follow accordingly.

(1) Sparsity. Recall that under the assumption of Theorem 1, β_0 lies in its feasible set conditional on \mathcal{E} . Since $\lambda_0 \leq \lambda_1$, by the definition of the constrained Dantzig selector, we have $|n^{-1}\mathbf{X}^T(\mathbf{y} - \mathbf{X}\hat{\beta})| \preceq \lambda_1$, where \preceq is understood as componentwise no larger than. Conditional on the event \mathcal{E} , substituting \mathbf{y} by $\mathbf{X}\beta_0 + \varepsilon$ and applying the triangle inequality yield

$$\|n^{-1}\mathbf{X}^T\mathbf{X}\delta\|_\infty \leq 2\lambda_1. \quad (11)$$

Furthermore, Lemma 3.1 in Candès and Tao (2007) still applies for δ as long as the uniform uncertainty principle condition (4) holds. Together with (10) and (11), by applying the same argument as in the proof of Theorem 1.1 in Candès and Tao (2007), we obtain an L_2 -estimation loss bound of $\|\delta\|_2 \leq (Cs)^{1/2}\lambda_1$ with $C = 4^2/(1 - \delta - \theta)^2$ some positive constant.

Since both the true regression coefficient vector β_0 and the constrained Dantzig selector $\hat{\beta}$ lie in the constrained parameter space \mathcal{B}_λ , the magnitude of any nonzero component in both β_0 and $\hat{\beta}$ is no smaller than λ . It follows that on the set of falsely discovered signs, the component of $\delta = \hat{\beta} - \beta_0$ will be no smaller than λ . Then making use of the obtained L_2 -estimation loss bound $\|\delta\|_2 \leq (Cs)^{1/2}\lambda_1$, it is immediate that the number of falsely discovered signs is bounded from above by $Cs(\lambda_1/\lambda)^2$. Under the assumption of $\lambda \geq C^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$, we have $Cs(\lambda_1/\lambda)^2 \leq s(1 + \lambda_1/\lambda_0)^{-2} \leq s$.

(2) L_2 -estimation loss. We further exploit the technical tool of Lemma 3.1 in Candès and Tao (2007) to analyze the behavior of the estimation error $\delta = \hat{\beta} - \beta_0$. Let δ'_1 be a subvector of δ_2 consisting of the s largest components in magnitude, $\delta_3 = (\delta_1^T, (\delta'_1)^T)^T$, and \mathbf{X}_3 a submatrix of \mathbf{X} consisting of columns corresponding to δ_3 . We emphasize that δ_3 covers all nonzero components of δ since the number of falsely discovered signs is upper bounded by s , as showed in the previous paragraph. Therefore, $\|\delta_3\|_q = \|\delta\|_q$ for all $q > 0$. In view of the uniform uncertainty principle condition (4), an application of Lemma 3.1 in Candès and Tao (2007) results in

$$\|\delta_3\|_2 \leq (1 - \delta)^{-1}\|n^{-1}\mathbf{X}_3^T\mathbf{X}\delta\|_2 + \theta(1 - \delta)^{-1}s^{-1/2}\|\delta_2\|_1. \quad (12)$$

On the other hand, from the basic inequality (10) it is easy to see that $s^{-1/2}\|\delta_2\|_1 \leq s^{-1/2}\|\delta_1\|_1 \leq \|\delta_3\|_2$. Substituting it into (12) leads to

$$\|\delta\|_2 = \|\delta_3\|_2 \leq (1 - \delta - \theta)^{-1}\|n^{-1}\mathbf{X}_3^T\mathbf{X}\delta\|_2 \quad (13)$$

Hence, to develop a bound for the L_2 -estimation loss it suffices to find an upper bound for $\|n^{-1}\mathbf{X}_3^T\mathbf{X}\delta\|_2$.

Denote by A_1 , A_2 , and A_3 the index sets of correctly selected variables, missed true variables, and falsely selected variables, respectively. Let $A_{23} = A_2 \cup A_3$. Then we can

obtain from the definition of the constrained Dantzig selector along with its thresholding feature that $|n^{-1}\mathbf{X}_{A_1}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})| \preceq \lambda_0$, $|n^{-1}\mathbf{X}_{A_2}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})| \preceq \lambda_1$, and $|n^{-1}\mathbf{X}_{A_3}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})| \preceq \lambda_0$. Conditional on \mathcal{E} , substituting \mathbf{y} by $\mathbf{X}\boldsymbol{\beta}_0 + \boldsymbol{\varepsilon}$ and applying the triangle inequality give

$$\|n^{-1}\mathbf{X}_{A_1}^T\mathbf{X}\boldsymbol{\delta}\|_\infty \leq 2\lambda_0 \quad \text{and} \quad \|n^{-1}\mathbf{X}_{A_{23}}^T\mathbf{X}\boldsymbol{\delta}\|_\infty \leq \lambda_0 + \lambda_1. \quad (14)$$

We now make use of the technical result on sparsity. Since A_{23} denotes the index set of false positives and false negatives, its cardinality is also bounded by $Cs(\lambda_1/\lambda)^2$ with probability at least $1 - O(n^{-c})$. Therefore, by (14) we have

$$\begin{aligned} \|n^{-1}\mathbf{X}_3^T\mathbf{X}\boldsymbol{\delta}\|_2^2 &\leq \|n^{-1}\mathbf{X}_{A_1}^T\mathbf{X}\boldsymbol{\delta}\|_2^2 + \|n^{-1}\mathbf{X}_{A_{23}}^T\mathbf{X}\boldsymbol{\delta}\|_2^2 \\ &\leq s\|n^{-1}\mathbf{X}_{A_1}^T\mathbf{X}\boldsymbol{\delta}\|_\infty^2 + Cs(\lambda_1/\lambda)^2\|n^{-1}\mathbf{X}_{A_{23}}^T\mathbf{X}\boldsymbol{\delta}\|_\infty^2 \\ &\leq 4s\lambda_0^2 + Cs(\lambda_1/\lambda)^2(\lambda_0 + \lambda_1)^2. \end{aligned}$$

Substituting this inequality into (13) yields

$$\|\boldsymbol{\delta}\|_2 \leq (1 - \delta - \theta)^{-1} \{4s\lambda_0^2 + Cs(\lambda_1/\lambda)^2(\lambda_0 + \lambda_1)^2\}^{1/2}.$$

Since we assume $\lambda \geq C^{1/2}(1 + \lambda_1/\lambda_0)\lambda_1$, it follows that $C(\lambda_1/\lambda)^2(\lambda_0 + \lambda_1)^2 \leq \lambda_0^2$. Therefore, we conclude that $\|\boldsymbol{\delta}\|_2 \leq (1 - \delta - \theta)^{-1}(5s\lambda_0^2)^{1/2}$.

(3) Other losses. Applying the basic inequality (10), we establish an upper bound for the L_1 -estimation loss

$$\begin{aligned} \|\boldsymbol{\delta}\|_1 &= \|\boldsymbol{\delta}_1\|_1 + \|\boldsymbol{\delta}_2\|_1 \leq 2\|\boldsymbol{\delta}_1\|_1 \leq 2s^{1/2}\|\boldsymbol{\delta}_1\|_2 \\ &\leq 2s^{1/2}\|\boldsymbol{\delta}\|_2 \leq 2(1 - \delta - \theta)^{-1}s(5\lambda_0^2)^{1/2}. \end{aligned}$$

For the L_∞ -estimation loss, we additionally assume that $\lambda > (1 - \delta - \theta)^{-1}(5s\lambda_0^2)$ which can lead to the sign consistency, $\text{sgn}(\hat{\boldsymbol{\beta}}) = \text{sgn}(\boldsymbol{\beta}_0)$, in view of the L_2 -estimation loss inequality above. Therefore, by the constrained Dantzig selector constraints we have $\|n^{-1}\mathbf{X}_1^T(\mathbf{y} - \mathbf{X}_1\hat{\boldsymbol{\beta}}_1)\|_\infty \leq \lambda_0$ and thus $\|n^{-1}\mathbf{X}_1^T(\boldsymbol{\varepsilon} - \mathbf{X}_1\boldsymbol{\delta}_1)\|_\infty \leq \lambda_0$. Then conditional on \mathcal{E} , it follows from the triangle inequality that $\|n^{-1}\mathbf{X}_1^T\mathbf{X}_1\boldsymbol{\delta}_1\|_\infty \leq 2\lambda_0$. Hence, $\|\boldsymbol{\delta}\|_\infty = \|\boldsymbol{\delta}_1\|_\infty \leq 2\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty\lambda_0$, which completes the proof.

Proof of Theorem 2

We continue to use the technical setup and notation introduced in the proof of Theorem 1. Results are parallel to those in Theorem 1 but presented in the asymptotic manner. The similarity lies in the rationale of the proof as well. The key element is to derive the sparsity and then construct an inequality for L_2 -estimation loss through the bridge $n^{-1}\|\mathbf{X}\boldsymbol{\delta}\|_2^2$. Inequalities for other types of losses are built upon this bound.

(1) Sparsity. Conditional on the event \mathcal{E} , we know that (10) and (11) still hold by the definition of the constrained Dantzig selector. Moreover, Condition 1 is similar to the restricted eigenvalue assumption $\text{RE}(s, m, 1)$ in Bickel et al. (2009), except that $\text{RE}(s, m, 1)$ assumes the inequality holds for any subset with size no larger than s to cover different possibilities of $\text{supp}(\boldsymbol{\beta}_0)$. Since we assume without loss of generality that $\text{supp}(\boldsymbol{\beta}_0) =$

$\{1, \dots, s\}$, an application of similar arguments as in the proof of Theorem 7.1 in Bickel et al. (2009) yields the L_2 oracle inequality $\|\boldsymbol{\delta}\|_2 \leq (C_m s)^{1/2} \lambda_1$ with C_m some positive constant dependent on m . Therefore, by the same arguments as in the proof of Theorem 1, it can be shown that the number of falsely discovered signs is bounded from above by $C_m s (\lambda_1/\lambda)^2$ and further by $s(1 + \lambda_1/\lambda_0)^{-2}$ since $\lambda \geq C_m^{1/2} (1 + \lambda_1/\lambda_0) \lambda_1$. Next we will go through the proof of Theorem 7.1 in Bickel et al. (2009) in a more cautious manner and make some improvements in some steps with the aid of the obtained bound on the number of false positives and false negatives.

(2) L_2 -estimation loss. By Condition 1, we have a lower bound for $n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2$. It is also natural to derive an upper bound for it and build an inequality related to the L_2 -estimation loss. It follows from (14) that

$$\begin{aligned} n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 &\leq \|n^{-1} \mathbf{X}_{A_1}^T \mathbf{X}\boldsymbol{\delta}\|_\infty \|\boldsymbol{\delta}_{A_1}\|_1 + \|n^{-1} \mathbf{X}_{A_{23}}^T \mathbf{X}\boldsymbol{\delta}\|_\infty \|\boldsymbol{\delta}_{A_{23}}\|_1 \\ &\leq 2\lambda_0 \|\boldsymbol{\delta}_{A_1}\|_1 + (\lambda_0 + \lambda_1) \|\boldsymbol{\delta}_{A_{23}}\|_1. \end{aligned} \quad (15)$$

Since the cardinality of A_{23} is bounded by $C_m s (\lambda_1/\lambda)^2$, applying the Cauchy-Schwarz inequality to (15) leads to $n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq 2\lambda_0 \|\boldsymbol{\delta}_{A_1}\|_1 + (\lambda_0 + \lambda_1) \|\boldsymbol{\delta}_{A_{23}}\|_1 \leq 2\lambda_0 s^{1/2} \|\boldsymbol{\delta}_{A_1}\|_2 + C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda \|\boldsymbol{\delta}_{A_{23}}\|_2$. This gives an upper bound for $n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2$. Combining this with Condition 1 results in

$$2^{-1} \kappa^2 (\|\boldsymbol{\delta}_{A_1}\|_2^2 + \|\boldsymbol{\delta}_{A_{23}}\|_2^2) \leq 2\lambda_0 s^{1/2} \|\boldsymbol{\delta}_{A_1}\|_2 + C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda \|\boldsymbol{\delta}_{A_{23}}\|_2. \quad (16)$$

Consider (16) in a two-dimensional space with respect to $\|\boldsymbol{\delta}_{A_1}\|_2$ and $\|\boldsymbol{\delta}_{A_{23}}\|_2$. Then the quadratic inequality (16) defines a circular area centered at $(2\kappa^{-2} \lambda_0 s^{1/2}, \kappa^{-2} C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda)$. The term $\|\boldsymbol{\delta}_{A_1}\|_2^2 + \|\boldsymbol{\delta}_{A_{23}}\|_2^2$ is nothing but the squared distance between the point in this circular area and the origin. One can easily identify the largest squared distance which is also the upper bound for the L_2 -estimation loss

$$\|\boldsymbol{\delta}\|_2^2 = \|\boldsymbol{\delta}_{A_1}\|_2^2 + \|\boldsymbol{\delta}_{A_{23}}\|_2^2 \leq 4\kappa^{-4} \{ (2\lambda_0 s^{1/2})^2 + [C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda]^2 \}.$$

With the assumption of $\lambda \geq C_m^{1/2} (1 + \lambda_1/\lambda_0) \lambda_1$, we can show that $\{C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda\}^2 \leq s\lambda_0^2$ and thus $\|\boldsymbol{\delta}\|_2 = O(\kappa^{-2} s^{1/2} \lambda_0)$. This bound has significance improvement with the factor $\log p$ reduced to $\log n$ in the ultra-high dimensional setting.

(3) Other losses. With the L_2 oracle inequality at hand, one can derive the L_1 oracle inequality $\|\boldsymbol{\delta}\|_1 \leq s^{1/2} \|\boldsymbol{\delta}\|_2 = O(\kappa^{-2} s \lambda_0)$ in a straightforward manner. For the prediction loss, it follows from (15) that

$$n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq 2\lambda_0 s^{1/2} \|\boldsymbol{\delta}_{A_1}\|_2 + C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda \|\boldsymbol{\delta}_{A_{23}}\|_2.$$

Consider the problem of minimizing $2\lambda_0 s^{1/2} \|\boldsymbol{\delta}_{A_1}\|_2 + C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda \|\boldsymbol{\delta}_{A_{23}}\|_2$ subject to (16). It is simply a two-dimensional linear optimization problem with a circular area as its feasible set. One can easily solve this minimization problem and obtain

$$n^{-1} \|\mathbf{X}\boldsymbol{\delta}\|_2^2 \leq 2\kappa^{-2} [(2\lambda_0 s^{1/2})^2 + \{C_m^{1/2} (\lambda_0 + \lambda_1) s^{1/2} \lambda_1 / \lambda\}^2] = O(\kappa^{-2} s \lambda_0^2).$$

Finally, for the L_∞ oracle inequality, it follows from similar arguments as in the proof of Theorem 1 that $\|\boldsymbol{\delta}\|_\infty \leq 2\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty\lambda_0 = O\{\|(n^{-1}\mathbf{X}_1^T\mathbf{X}_1)^{-1}\|_\infty\lambda_0\}$, which concludes the proof.

Proof of Theorem 3

Denote by $\hat{\boldsymbol{\beta}}_{\text{global}}$ the global minimizer of (3). Under the conditions of Theorem 2, $\hat{\boldsymbol{\beta}}_{\text{global}}$ enjoys the same oracle inequalities and properties as in Theorem 2 conditional on the event defined in (7). In particular, we have $FS(\hat{\boldsymbol{\beta}}_{\text{global}}) = O(s)$. It follows that

$$\|\hat{\boldsymbol{\beta}}_{\text{global}}\|_0 \leq \|\hat{\boldsymbol{\beta}}_0\|_0 + FS(\hat{\boldsymbol{\beta}}_{\text{global}}) = O(s).$$

Let $\hat{\boldsymbol{\beta}}$ be a computable local minimizer of (3) produced by any algorithm satisfying $\|\hat{\boldsymbol{\beta}}\|_0 \leq c_2s$. Denote by $A = \text{supp}(\hat{\boldsymbol{\beta}}) \cup \text{supp}(\hat{\boldsymbol{\beta}}_{\text{global}})$. Then we have

$$|A| \leq \|\hat{\boldsymbol{\beta}}\|_0 + \|\hat{\boldsymbol{\beta}}_{\text{global}}\|_0 = O(s) \leq c_3s$$

for some large enough positive constant c_3 .

We next analyze the difference between the two estimators, that is, $\boldsymbol{\delta} = \hat{\boldsymbol{\beta}} - \hat{\boldsymbol{\beta}}_{\text{global}}$. Let \mathbf{X}_A be a submatrix of \mathbf{X} consisting of columns in A and $\boldsymbol{\delta}_A$ a subvector of $\boldsymbol{\delta}$ consisting of components in A . Since $\|\mathbf{X}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})\|_\infty = O(\lambda_1)$ and $\|\mathbf{X}^T(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}_{\text{global}})\|_\infty = O(\lambda_1)$, we can show that

$$\|n^{-1}\mathbf{X}_A^T\mathbf{X}_A\boldsymbol{\delta}_A\|_2 \leq |A|^{1/2}\|n^{-1}\mathbf{X}_A^T\mathbf{X}_A\boldsymbol{\delta}_A\|_\infty = O(s^{1/2}\lambda_1).$$

By the assumption that $\min_{\|\boldsymbol{\delta}\|_2=1, \|\boldsymbol{\delta}\|_0 \leq c_3s} n^{-1/2}\|\mathbf{X}\boldsymbol{\delta}\|_2 \geq \kappa_0$, the smallest singular value of $n^{-1/2}\mathbf{X}_A$ is bounded from below by κ_0 . Thus we have $\|\boldsymbol{\delta}\|_2 = \|\boldsymbol{\delta}_A\|_2 = O(s^{1/2}\lambda_1)$. Together with the thresholding feature of the constrained Dantzig selector, it follows that the number of different indices between $\text{supp}(\hat{\boldsymbol{\beta}})$ and $\text{supp}(\hat{\boldsymbol{\beta}}_{\text{global}})$ is bounded by $O\{s(\lambda_1/\lambda)^2\}$. This sparsity property is essential for our proof and similar sparsity results can be found in Theorem 2.

Based on the aforementioned sparsity property and by the facts that $\|\mathbf{X}_{\tilde{A}_1}^T\mathbf{X}\boldsymbol{\delta}\|_\infty \leq 2\lambda_0$ and $\|\mathbf{X}_{\tilde{A}_{23}}^T\mathbf{X}\boldsymbol{\delta}\|_\infty \leq \lambda_0 + \lambda_1$ with $\tilde{A}_1 = \text{supp}(\hat{\boldsymbol{\beta}}) \cap \text{supp}(\hat{\boldsymbol{\beta}}_{\text{global}})$ and $\tilde{A}_{23} = [\text{supp}(\hat{\boldsymbol{\beta}}) \setminus \text{supp}(\hat{\boldsymbol{\beta}}_{\text{global}})] \cup [\text{supp}(\hat{\boldsymbol{\beta}}_{\text{global}}) \setminus \text{supp}(\hat{\boldsymbol{\beta}})]$, the same arguments as in the proof of Theorem 2 apply to show that $\|\boldsymbol{\delta}\|_2 = O(\kappa^{-2}s^{1/2}\lambda_0)$. It is clear that this bound is of the same order as that for the difference between $\hat{\boldsymbol{\beta}}_{\text{global}}$ and $\boldsymbol{\beta}_0$ in Theorem 2. Thus $\hat{\boldsymbol{\beta}}$ enjoys the same asymptotic bound on the L_2 -estimation loss. Similarly, the asymptotic bounds for the other losses in Theorem 2 also apply to $\hat{\boldsymbol{\beta}}$, since those inequalities are rooted on the bounds for the sparsity and L_2 -estimation loss. This completes the proof.

Antoniadis, A., Fryzlewicz, P. and Letue, F. (2010). The Dantzig selector in Cox's proportional hazards model. *Scand. J. Statist.* **37**, 531–552.

- Bickel, P. J., Ritov, Y. and Tsybakov, A. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **37**, 1705–1732.
- Candès, E. J. and Tao, T. (2005). Decoding by linear programming. *IEEE Trans. Info. Theory* **51**, 4203–4215.
- Candès, E. J. and Tao, T. (2007). The Dantzig selector: statistical estimation when p is much larger than n (with Discussion). *Ann. Statist.* **35**, 2313–2351.
- Candès, E. J., Wakin, M. B. and Boyd, S. (2008). Enhancing sparsity by reweighted l_1 minimization. *J. Fourier Anal. Appl.* **14**, 877–905.
- Chen, K., Dong, H. and Chan, K. S. (2013). Reduced rank regression via adaptive nuclear norm penalization. *Biometrika* **100**, 901–920.
- Donoho, D. L. (2006). Compressed sensing. *IEEE Trans. Info. Theory* **52**, 1289–1306.
- Dudley, R. M. (1999). *Uniform Central Limit Theorems*. Cambridge University Press.
- Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96**, 1348–1360.
- Fan, J. and Lv, J. (2011). Nonconcave penalized likelihood with NP-dimensionality. *IEEE Trans. Info. Theory* **57**, 5467–5484.
- Fan, J. and Peng, H. (2004). Nonconcave penalized likelihood with diverging number of parameters. *Ann. Statist.* **32**, 928–961.
- Fan, Y. and Lv, J. (2013). Asymptotic equivalence of regularization methods in thresholded parameter space. *J. Am. Statist. Assoc.* **108**, 1044–1061.
- Fan, Y. and Lv, J. (2014). Asymptotic properties for combined L_1 and concave regularization. *Biometrika* **101**, 57–70.
- Fan, Y. and Tang, C. (2013). Tuning parameter selection in high dimensional penalized likelihood. *J. R. Statist. Soc. Ser. B* **75**, 531–552.
- van de Geer, S. (2008). High-dimensional generalized linear models and the Lasso. *Ann. Statist.* **36**, 614–645.
- van de Geer, S., Bühlmann, P. and Zhou S. (2011). The adaptive and the thresholded Lasso for potentially misspecified models (and a lower bound for the Lasso). *Electronic Journal of Statistics* **5**, 688–749.
- Huang, J., Ma, S. and Zhang, C.-H. (2008). Adaptive Lasso for sparse high-dimensional regression models. *Statistica Sinica* **18**, 1603–1618.
- Lan, H., Chen, M., Flowers, J., Yandell, B., Stapleton, D., Mata, C., Mui, E., Flowers, M., Schueler, K., Manly, K., Williams, R., Kendzioriski, K. and Attie, A. D. (2006). Combined expression trait correlations and expression quantitative trait locus mapping. *PLoS Genet.* **2**, e6.

- Lv, J. and Fan, Y. (2009). A unified approach to model selection and sparse recovery using regularized least squares. *Ann. Statist.* **37**, 3498–3528.
- Negahban, S., Ravikumar, P., Wainwright, M. and Yu, B. (2012). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science* **27**, 538–557.
- Raskutti, G., Wainwright, M. J. and Yu, B. (2011). Minimax rates of convergence for high-dimensional regression under ℓ_q -ball sparsity. *IEEE Trans. Inform. Th.* **57**, 6976–6994.
- Scheetz, T., Kim, K., Swiderski, R., Philp, A., Braun, T., Knudtson, K., Dorrance, A., DiBona, G., Huang, J., Casavant, T. and et al. (2006). Regulation of gene expression in the mammalian eye and its relevance to eye disease. *Proc. Natl. Acad. Sci. USA* **103**, 14429–14434.
- Song, Q. and Liang, F. (2015). High dimensional variable selection with reciprocal L_1 -regularization. *J. Am. Statist. Assoc.*, in press.
- Städler, N., Bühlmann, P. and van de Geer, S. (2010). ℓ_1 -penalization in mixture regression models (with discussion). *Test* **19**, 209–285.
- Tibshirani, R. J. (1996). Regression shrinkage and selection via the Lasso. *J. R. Statist. Soc. Ser. B* **58**, 267–288.
- Wainwright, M. J. (2009) Sharp thresholds for noisy and high-dimensional recovery of sparsity using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Th.* **55**, 2183–2202.
- Xue, L. and Zou, H. (2011). Sure independence screening and compressed random sensing. *Biometrika* **98**, 371–380.
- Zhang, C.-H. (2010) Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38**, 894–942.
- Zhang, T. (2011) Adaptive forward-backward greedy algorithm for learning sparse representations. *IEEE Trans. Inform. Th.* **57**, 4689–4708.
- Zhao, P. and Yu, B. (2006). On model selection consistency of Lasso. *J. Machine Learning Research* **7**, 2541–2567.
- Zheng, Z., Fan, Y. and Lv, J. (2014). High-dimensional thresholded regression and shrinkage effect. *J. R. Statist. Soc. Ser. B* **76**, 627–649.
- Zhou, T. and Tao, D. (2013). Double shrinking sparse dimension reduction. *IEEE Transactions on Image Processing* **22**, 244–257.
- Zou, H. (2006). The adaptive Lasso and its oracle properties. *J. Am. Statist. Assoc.* **101**, 1418–1429.
- Zou, H. and Hastie, T. J. (2005). Regularization and variable selection via the elastic net. *J. R. Statist. Soc. Ser. B* **67**, 301–320.